Network Security

NetDetector – Network Intrusion Forensic System (NIFS)

White paper

www.niksun.com

NIKSUN, NetVCR, NetDetector, NetX, NetXperts, NetReporter, and NSS are either registered trademarks or trademarks of NIKSUN, Inc. in the United States and/or other countries.

Ethernet is a trademark of Xerox Corp.
Netscape Communicator is a trademark of Netscape Communications Corporation.
Internet Explorer is a trademark of Microsoft Corporation.

Other product and company names mentioned herein may be the trademarks of their respective owners.

Copyright © 2003 NIKSUN, Inc.

This publication is protected by International Copyright Law. No part of this publication may be reproduced, stored in a retrieval system, translated, transcribed, or transmitted in any form, or by any means manual, electric, electronic, electromagnetic, mechanical, chemical, optical, or otherwise, without prior written permission from NIKSUN, Inc. NIKSUN makes no warranty of any kind with respect to this material and disclaims any implied warranty of merchantability or fitness for a particular purpose.

NIKSUN, Inc.
1100 Cornwall Road
Monmouth Junction, NJ 08852
USA
Telephone: (732) 821-5000
Fax: (732) 821-6000
Customer Support: (888) 821-2003
E-mail:info@NIKSUN.com

Table of Contents

1. Introduction 1

2. NetDetector – Network Intrusion Forensic System (NIFS) 2

Intrusion detection 2
Forensic analysis 2

3. How to Use NetDetector: Sample Scenarios 4

Scenario 1: Investigating for A Critical Intrusion Alert and Reduce/Eliminate False Positives 4
NetDetector 4
IDS 4
Scenario 2: Investigating the Activities of a Particular Host 5
Scenario 3: New Attack, Signature Update, and Verification 6
NetDetector 6
IDS 6

4. Conclusion 8

1. Introduction

Every system connected to the Internet is a potential candidate for a malicious attack. Even the most secure Internet sites, corporations, and government agencies have experienced security violations and network attacks. To deal with those risks, organizations are deploying network infrastructure devices such as firewalls and intrusion detection systems (IDSs) to increase the security of their networks. However, the notion of an impenetrable network is equivalent to the notion of an unsinkable ship, and one need not be reminded of the lessons learned from the Titanic!

Due to the complexity of network attacks and the extent of the damage they cause, organizations are spending considerably more time and resources in recovering from security incidents than in the past. Viruses and worms, such as Nimda and Code Red, have caused major disruptions of service on corporate networks and operation. It has become critical that organizations deploy technologies that help reduce the time and cost of recovering from security incidents.

In order to protect computer networks against unauthorized access, intrusion attempts, and network attacks, best practices dictate deploying multiple layers of devices and technologies to ensure that networks are secure and can recover rapidly from security breaches. NIKSUN recommends a three-tiered security approach:

TIER	DESCRIPTION
AVOIDANCE	Apply protective mechanisms such as firewalls, VPNs, encryption, and authentication technologies to prevent unauthorized access and intrusion attacks from occurring in the first place.
DETECTION	Try to detect intrusions and attempted intrusions that do occur by utilizing such technologies as host-based and network-based IDSs and virus scanners.
INVESTIGATION	Collect necessary information to enable users to recover from the damage if security breaches do occur, take steps to prevent breaches from happening again, and substantiate policy offenses.

Clearly, the IDS plays its main role in the *Detection Tier* of the NIKSUN security taxonomy. In particular, the IDS scans for known attack patterns and generates alarms when those patterns are detected in the network traffic or in host logs.

With its threshold-based anomaly alarms and its support of signature-based intrusion alarms (through its integration of Snort), NIKSUN's NetDetector plays a role within the Detection Tier. However, NetDetector also plays a critical role within the *Investigation Tier*. By recording all network traffic and providing long-term storage of raw packets and statistics, NetDetector allows security administrators to go back in time, recreate all network events, and conduct a complete forensic investigation of network intrusions.

2. NetDetector – Network Intrusion Forensic System (NIFS)

Intrusion detection

Network intrusion detection technologies provide security administrators with the ability to detect intrusion attempts. Most network-based IDSs (e.g., ISS RealSecure, Cisco Secure IDS, and Snort) include a signature database to detect a large number of ***known*** attacks. Security administrators will typically monitor the alerts and investigate the most critical ones as they can help identify major security breaches. They also periodically update the IDS signature database since new viruses and attacks are introduced continuously.

However, when a new virus or attack is introduced on the network, the database of known signatures would be unable to detect it till the new signature is identified and downloaded. This implies that networks are always vulnerable to new attacks. In addition, the time delay between a new attack being publicized and signatures being updated creates an exploitable window of opportunity for enterprising hackers. Furthermore, once the network security is compromised, it is very difficult to determine which system has been infected, what kind of damage has been caused, and if any intellectual property has been compromised. For example, when the Nimda virus was introduced, before it was identified the IDS signature database was updated, many corporate networks were infected, operations interrupted, and millions of dollars of revenues were lost. Though an IDS generates alarms against potential security violations, it does not provide much, if any, mechanism for the security administrators to conduct an investigation. Furthermore, it does not provide the security administrators any tools to verify if the network has been compromised in the time frame that a new virus or attack is introduced and the known signature database is updated. Using NetDetector as the IDS can eliminate these drawbacks.

False positive alarms are another major drawback of an IDS. Gartner Group declares, “By the end of 2003, 90% of IDS deployments will fail when false positives are not reduced by 90%”. By applying NetDetector’s signature database repeatedly against its recorded static dataset, the security administrators can fine-tune NetDetector (Snort) signature database to reduce and eliminate false positive alarms and verify custom rules are accurately created to capture potential intrusions

Forensic analysis

NetDetector is the market leader in the network forensic and analysis segment. It enables security administrators to investigate network security incidents caused by both external intrusions and internal security breaches. NetDetector provides the security administrators an integrated network security solution that is capable of detecting both signature based and anomaly based intrusions and equipped with powerful investigation capabilities. NetDetector’s post-event investigation capabilities can provide the causal link between distant, incongruent, and incomplete events generated by other IDS.

In many cases, hacking a computer system involves multiple coordinated actions, not all of which generate IDS alerts. For instance, a hacker could compromise a host using a known buffer-overflow attack, install some root-kits, and then initiate attacks from that host. The security administrator might get IDS alerts that do not provide enough data to determine the extent of the damage. NetDetector provides all the pieces of the puzzle for the end-to-end analysis.

For those alerts generated by an IDS (both NetDetector and others), the security administrator can utilize the data recorded by NetDetector to conduct a security investigation to verify if the IDS alert is a false alarm or an actual security breach. Furthermore, since NetDetector continuously records all network traffic, when the network security is compromised NetDetector can provide the security administrators a complete audit trail to determine which host has been attacked, if intellectual property has been compromised, if any host has been "hi-jacked", and how it occurred. Using this information, security administrators can more rapidly recover from damages to the network and repair security loopholes.

3. How to Use NetDetector: Sample Scenarios

Utilizing alert data from an IDS alone (e.g., timestamps, IP addresses, port numbers, protocols) for investigation is potentially a very time consuming process and very likely to yield no positive results due to lack of packet data. Though using NetDetector's large disk storage for post-event analysis has been likened to finding the proverbial "needle in a haystack", when used in conjunction with alert data from an IDS (both NetDetector itself and others) a potentially very time consuming investigation process could be shortened significantly and greatly simplified. NetDetector captures all the traffic to enable the full investigation and its intrusion detection feature provides the pointers to start the investigation.

Following are some scenarios that illustrate the process of taking advantage of NetDetector's intrusion detection and forensic analysis capabilities.

Scenario 1: Investigating for A Critical Intrusion Alert and Reduce/Eliminate False Positives

Assume that NetDetector (or an IDS) detected a particular TCP-based signature and raised a critical alert. The alert data from NetDetector (or the IDS) provides information required to identify the flow that contains the signature. To perform the event investigation the security administrator can do the following:

NetDetector

1. Double click the specific alert of interests from the Event Viewer screen to bring up the Event Details screen
2. Click the View Analysis button of the Event Details to bring up the Analysis screen. Note, the Analysis screen displays statistic data and information of the time interval around the event timestamp (e.g., +/- 3 minutes) and is filtered on the client and server IP addresses
3. Drill down into the TCP protocol
4. List all the TCP sessions between the IP addresses
5. If required, filter on the client and server port numbers
6. Display the list of TCP sessions, reassemble the data, and view the packet dump.
7. Archive the data if required

IDS

1. Extract the timestamp, host addresses, and port numbers from the Event Detail window of the IDS
2. Select a time interval around the event timestamp (e.g., +/- 5 minutes) and select the IP protocol on the NetDetector home screen.
3. Filter on the client and server IP addresses
4. Verify the IP protocols used between the two IP addresses

5. Drill down into the TCP protocol
6. List all the TCP sessions between the IP addresses
7. If required, filter on the client and server port numbers
8. Display the list of TCP sessions, reassemble the data, and view the packet dump.
9. Archive the data if required

The incident investigation capabilities provided by NetDetector enable the security analyst to determine whether the network attack was real or a false positive and whether the attack was successful (in breaching the network integrity). The security analyst can also evaluate the extent of the damage on the compromised host and investigate any other action taken by the intruder.

When NetDetector is used as an IDS and this alert is a false positive, the security analyst can repeatedly apply the recorded dataset against NetDetector (Snort) signature database to fine-tune its database to reduce and eliminate false positives.

A simple example, assume the alert was on the well-known PHF web server attack (CVE-1999-0067). This attack is often identified by the presence of "/cgi-bin/phf?" in an HTTP POST request, and can allow arbitrary commands to be executed on the web server. When viewing the packets in Step 6 (or Step 8 for an IDS) above, it will be evident if the activity was an attack or benign (a false positive), such as a user entering "/cgi-bin/phf?" in a search engine. If it turns out to be an attack, the security analyst can then reconstruct the session to see exactly what arbitrary commands were executed and what information the hacker has accessed. They can then use the information for further exploits:

- If the hacker simply replaced some HTML files to deface the website, the analysis that NetDetector provides would tell the security analyst exactly which files were replaced so that they can be restored.

- If the hacker installed a backdoor to exploit at a later time, the exact name and location of the installed files would be viewable so that they could be immediately removed.

- If the hacker accessed a database of sensitive information (e.g., a customer database with credit card numbers or confidential documents), the exact information could be determined and corrective actions initiated (e.g., notifying the credit card holders, or flagging or changing their account numbers).

In all of these cases, the investigation and recovery could be done in hours by utilizing NetDetector rather than in weeks or months. And clearly, time is of the essence for recovery from security breaches to minimize revenue loss.

Scenario 2: Investigating the Activities of a Particular Host

IDS technologies are very effective at detecting known attacks. There are instances where particular IP addresses are involved in many security events. For example, NetDetector (or an IDS) could detect that an offending host conducted port-scans on five hosts and caused buffer-overflows in two hosts. This is clearly suspicious, and the security administrator should investigate all the activities associated with that offending host (IP address).

It is clear that determining the IP addresses to investigate does not depend on the existence of one particular event but on the existence of various events that make the IP address suspicious. The security administrators might want to investigate if the host is a victim of various attacks or if it has been initiating attacks. Using NetDetector to perform the event investigation, the security administrator can do the following:

1. Filter on the IP address
2. Verify the IP protocols used
3. Drill down into the TCP protocol
4. List all the TCP sessions
5. If required, filter on the port numbers
6. Display the list of TCP sessions, reassemble the data, and view the packet dump
7. Archive the data if required

NetDetector empowers the security analyst to determine whether the host is initiating an attack or has been "hi-jacked" (and is acting as a zombie). If the host has been "hi-jacked", the security analyst can also evaluate the extent of the damage on the compromised host and investigate if the intruder has taken any other actions. NetDetector's large storage and data export ability allow the potential to investigate all the activity over the course of days, weeks, or even months. Perhaps the initial compromise to install a backdoor happened last week and the launch of an attack from the compromised host happened today, only NetDetector would allow a security analyst to piece it all together since data is stored in its huge data storage.

Scenario 3: New Attack, Signature Update, and Verification

Malicious hackers are an industrious and prolific bunch. New exploits are being developed all the time, and invariably, one will slip through network perimeters without detection. The unfortunate network to get hit first is usually subjected to a detailed forensics investigation (another point where having a NetDetector is beneficial!) and often a new signature is developed. It then takes time for IDS vendors to update their signatures, notify their customers, and then for end users to actually install the updates. What happens between the times that the new attack is first discovered and IDSs are all updated with the new signature? You can bet that the opportunistic hackers are busy targeting as many networks as they can with the new attack before they can be detected! Has your network been one of them? If you had a NetDetector, it would be relatively easy to tell:

NetDetector

1. New attack is announced at time t0.
2. Update your NetDetector (Snort) signature database at time t1.
3. Apply NetDetector (Snort) signature database against the data you have recorded between t0 and t1 (by changing the Monitoring Interface field of the Configure Snort IDS screen).
4. If the new signature is triggered, then you know you were hit with the attack. From here you can take corrective actions.

IDS

1. New attack is announced at time t0.
2. Update your IDS signatures at time t1.

3. With a NetDetector, analyze the data you have recorded between t0 and t1.
4. Replay the data between t0 and t1 onto a non-production network segment where the updated IDS is monitoring.
5. If the new signature is triggered, then you know you were hit with the attack. From here you can take corrective actions.

As an alternative, there are some cases where you can determine if you have been hit with the new attack even before updating the IDS signatures. If, for example, the attack is characterized by some well-known string of characters, you can use NetDetector's string search capabilities to scan for the attack over all recorded data.

Same process can also be used to verify if a custom rule created by the security administrator was constructed accurately.

4. Conclusion

A network security architecture includes NIKSUN NetDetectors with or without IDSs provides a powerful intrusion detection and forensic analysis capability, beyond what an IDS could do alone. The IDS (and NetDetector) provides powerful signature detection and alerting capabilities and has a large database of signatures for known security vulnerabilities. However, the IDS alone lacks the ability to provide long term and large-scale storage of raw packet data. Furthermore, it does not archive any data it fails to (or was not configured to) detect. NetDetector, on the other hand, provides continuous recording and a large-scale storage, thus enabling invaluable post-event analysis of all network events. In addition, NetDetector includes network anomaly detection based on statistical traffic analysis. NetDetector (or a solution that combines NetDetector and an IDS) would provide the following benefits:

- **Security incident identification and forensic analysis:** NetDetector's (or an IDS') signature detection capabilities would complement the anomaly detection of NetDetector. Given valuable alert information from NetDetector (or an IDS), security administrators can pinpoint and analyze the details of the incident using NetDetector. The large-scale storage and analysis capabilities of NetDetector would complement the IDS detection. NetDetector users are able to quickly examine the intrusion, determine the extent of the damage, and implement a complete recovery. Furthermore, they are able to export the data to a third party system for longer-term storage or analysis.

- **False positive handling**: The analysis capabilities of NetDetector can be used to drill down on data to quickly determine false positives alerts, hence saving investigation time for security analysts. Data can then be used to fine-tune NetDetector (or an IDS) signature (alert).

- **False negative investigation:** Since NetDetector continuously records data, any events that may not have generated alerts (in the IDS) will still be available for analysis. In addition, various IDS evasion techniques that may bypass the IDS (or Snort) will still be recorded by NetDetector.

- **Signature tuning and creation:** The forensic analysis of false negatives will provide valuable clues for fine tuning or adding new intrusion detection signatures or custom rules. Data recorded by NetDetector (or NetDetector's replay capability) can be used to test newly updated signatures in Snort (or the IDS), to determine if the new attacks are present in past data, and to verify the correctness of the new signatures or custom rules.

- **Signature Detection in data from WAN interfaces:** NetDetector supports WAN interfaces. It is also able to replay data captured on WAN interfaces onto an Ethernet link. Therefore it can be used to stream such data into IDSs (connected to a LAN) that in turn performs the IDS signature detection and alerting, thus expanding the capabilities of the IDS to analyze the WAN interfaces.